arXiv:1605.08963v3 [stat.ME] 3 Jun 2016

Variable selection in seemingly unrelated regressions with random predictors

David Puelz, P. Richard Hahn and Carlos M. Carvalho

David Puelz
e-mail: david.puelz@utexas.edu

P. Richard Hahn
e-mail: richard.hahn@chicagobooth.edu

Carlos M. Carvalho
e-mail: carlos.carvalho@mccombs.utexas.edu

Abstract: This paper considers linear model selection when the response is vector-valued and the predictors are randomly observed. We propose a new approach that decouples statistical inference from the selection step in a "post-inference model summarization" strategy. We study the impact of predictor uncertainty on the model selection procedure. The method is demonstrated through an application to asset pricing.

1. Introduction and overview

This paper develops a method for parsimoniously summarizing the shared dependence of many individual response variables upon a common set of predictor variables drawn at random. The focus is on multivariate Gaussian linear models where an analyst wants to find, among p available predictors X, a subset which work well for predicting $q > 1$ response variables Y. The multivariate normal linear model assumes that a set of responses $\{Y_j\}_{j=1}^q$ are linearly related to a shared set of covariates $\{X_i\}_{i=1}^p$ via

$$Y_j = \beta_{j1}X_1 + \cdots + \beta_{jp}X_p + \epsilon_j, \qquad \boldsymbol{\epsilon} \sim \mathrm{N}(0, \Psi), \tag{1}$$

where Ψ is a non-diagonal covariance matrix.

Bayesian variable selection in (single-response) linear models is the subject of a vast literature, from prior specification on parameters (Bayarri et al., 2012) and models (Scott and Berger, 2006) to efficient search strategies over the model space (George and McCulloch, 1993; Hans, Dobra and West, 2007). For a more complete set of references we refer the reader to the reviews of Clyde and George (2004) and Hahn and Carvalho (2015). By comparison, variable selection has not been widely studied in concurrent regression models, perhaps because it is natural simply to apply existing variable selection methods to each univariate regression individually. Indeed, such joint regression models go by the name "seemingly unrelated regressions" (SUR) in the Bayesian econometrics literature, reflecting the fact that the regression coefficients from each of the separate regressions can be obtained in isolation from one another (i.e., conducting estimation as if Ψ were diagonal). However, allowing non-diagonal Ψ can lead to more efficient estimation (Zellner, 1962) and can similarly impact variable selection (Brown, Vannucci and Fearn, 1998; Wang, 2010).

This paper differs from Brown, Vannucci and Fearn (1998) and Wang (2010) in that we focus on the case where the predictor variables (the regressors, or covariates) are treated as random as opposed to fixed.Our goal will be to summarize codependence among multiple responses in *subsequent* periods, making the uncertainty in future realizations highly central to our selection objective. This approach

is natural in many contexts (e.g., macroeconomic models) where the purpose of selection is inherently forward-looking. To our knowledge, no existing variable selection methods are suitable in this context. The new approach is based on the sparse summary perspective outlined in Hahn and Carvalho (2015), which applies Bayesian decision theory to summarize complex posterior distributions. By using a utility function that explicitly rewards sparse summaries, a high dimensional posterior distribution is collapsed into a more interpretable sequence of sparse point summaries.

A related approach to variable selection in multivariate Gaussian models is the Gaussian graphical model framework (Jones et al., 2005; Dobra et al., 2004; Wang and West, 2009). In that approach, the full conditional distributions are represented in terms of a sparse $(p+q)$-by-$(p+q)$ precision matrix. By contrast, we partition the model into response and predictor variable blocks, leading to a distinct selection criterion that narrowly considers the p-by-q covariance between Y to X.

1.1. Methods overview

Posterior summary variable selection consists of three phases: *model specification and fitting*, *utility specification*, and *graphical summary*. Each of these steps is outlined below. Additional details of the implementation are described in Section 2 and the Appendix.

Step 1: Model specification and fitting

The statistical model may be described compositionally as $p(Y, X) = p(Y|X)p(X)$. For $(Y, X) \sim \mathrm{N}(\mu, \Sigma)$, the regression model (1) implies Σ has the following block structure:

$$\Sigma = \left[\begin{array}{c|c} \boldsymbol{\beta}^T\Sigma_x\boldsymbol{\beta} + \Psi & (\Sigma_x\boldsymbol{\beta})^T \\ \hline \Sigma_x\boldsymbol{\beta} & \Sigma_x \end{array}\right]. \tag{2}$$

We denote the unknown parameters for the full joint model as $\Theta = \{\mu_x, \mu_y, \Sigma_x, \boldsymbol{\beta}, \Psi\}$ where $\mu = (\mu_y^T, \mu_x^T)^T$ and $\Sigma_x = \mathrm{cov}(X)$.

For a given prior choice $p(\Theta)$, posterior samples of all model parameters are computed by routine Monte Carlo methods, primarily Gibbs sampling. Details of the specific modeling choices and associated posterior sampling strategies are described in the Appendix.

A notable feature of our approach is that *steps 2* (and *3*) will be unaffected by modeling choices made in *step 1* except insofar as they lead to different posterior distributions $p(\Theta|\mathbf{Y}, \mathbf{X})$. In short, *step 1* is "obtain a posterior distribution"; posterior samples then become inputs to *step 2*.

Step 2: Utility specification

For our utility function we use the log-density of the regression $p(Y|X)$ above. It is convenient to work in terms of negative utility, or loss:

$$\mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}) = \frac{1}{2}(\tilde{Y} - \boldsymbol{\gamma}\tilde{X})^T\Omega(\tilde{Y} - \boldsymbol{\gamma}\tilde{X}), \tag{3}$$

where $\Omega = \Psi^{-1}$. Note that this log-density is being used in a descriptive capacity, not an inferential one; that is, all posterior inferences are based on the posterior distribution from *step 1*. The "action" $\boldsymbol{\gamma}$ is regarded as a point estimate of the

regression parameters $\boldsymbol{\beta}$, which would be a good fit to *future* data $(\tilde{Y}, \tilde{X})$ drawn from the same model as the observed data.

Taking expectations over the posterior distribution of all unknowns

$$p(\tilde{Y}, \tilde{X}, \Theta|\mathbf{Y}, \mathbf{X}) = p(\tilde{Y}|\tilde{X}, \Theta)p(\tilde{X}|\Theta)p(\Theta|\mathbf{Y}, \mathbf{X}), \tag{4}$$

yields expected loss

$$\mathcal{L}(\boldsymbol{\gamma}) \equiv \mathbb{E}[\mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma})] = \text{tr}[M\boldsymbol{\gamma}S\boldsymbol{\gamma}^T] - 2\text{tr}[A\boldsymbol{\gamma}^T] + \text{constant}, \tag{5}$$

where $A = \mathbb{E}[\Omega\tilde{Y}\tilde{X}^T]$, $S = \mathbb{E}[\tilde{X}\tilde{X}^T] = \overline{\Sigma_x}$, and $M = \overline{\Omega}$, the overlines denote posterior means, and the final term is a constant with respect to $\boldsymbol{\gamma}$.

Finally, we add an explicit penalty, reflecting our preference for sparse summaries:

$$\mathcal{L}_\lambda(\boldsymbol{\gamma}) \equiv \text{tr}[M\boldsymbol{\gamma}S\boldsymbol{\gamma}^T] - 2\text{tr}[A\boldsymbol{\gamma}^T] + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0, \tag{6}$$

where $\left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0$ counts the number of non-zero elements in $\boldsymbol{\gamma}$. In practice, we will use an approximation to this utility based on the ℓ_1 penalty; optimal actions under this approximation will still be sparse.

Step 3: Graphical summary

Traditional applications of Bayesian decision theory derive *point-estimates* by minimizing expected loss for certain loss functions. The present goal is not an *estimator* per se, but a parsimonious summary of information contained in a complicated, high dimensional posterior distribution. This distinction is worth emphasizing because we have not one, but rather a continuum of loss functions, indexed by the penalty parameter λ. This class of loss functions can be used to organize the posterior distribution as follows.

Using available convex optimization techniques, expression (6) can be optimized efficiently for a range of λ values simultaneously. Posterior graphical summaries consist of two components. First, graphs depicting which response variables have non-zero $\boldsymbol{\gamma}^*_\lambda$ coefficients on which predictor variables can be produced for any given λ. Second, posterior distributions of the quantity

$$\Delta_\lambda = \mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}^*_\lambda) - \mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}^*) \tag{7}$$

can be used to gauge the impact λ has on the descriptive capacity of γ^*_λ. Here, $\boldsymbol{\gamma}^* = \gamma^*_{\lambda=0}$ is the unpenalized optimal solution to the minimization of loss (6).

2. Posterior summary variable selection

The statistical model is given in equations (1) and (2); prior specification and model fitting details can be found in the Appendix. Alternatively, the models described in Brown, Vannucci and Fearn (1998) or Wang (2010) could be used. In this section, we flesh out the details of *steps 2* and *3*, which represent the main contributions of this paper.

2.1. Deriving the sparsifying expected utility function

Define the optimal posterior summary as the $\boldsymbol{\gamma}^*$ minimizing some expected loss $\mathcal{L}_\lambda(\boldsymbol{\gamma}) = \mathbb{E}[\mathcal{L}_\lambda(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma})]$. Here, the expectation is taken over the joint posterior predictive and posterior distribution: $p(\tilde{Y}, \tilde{X}, \Theta \mid \mathbf{Y}, \mathbf{X})$.

As described in the previous section, our loss takes the form of a penalized log conditional distribution:

$$\mathcal{L}_\lambda(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}) \equiv \frac{1}{2}(\tilde{Y} - \boldsymbol{\gamma}\tilde{X})^T \Omega (\tilde{Y} - \boldsymbol{\gamma}\tilde{X}) + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_0, \tag{8}$$

where $\Omega = \Psi^{-1}$, $\left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_0 = \sum_j \mathbb{1}\left(\mathbf{vec}(\boldsymbol{\gamma}) \neq 0\right)$, and $\mathbf{vec}(\boldsymbol{\gamma})$ is the vectorization of the action matrix $\boldsymbol{\gamma}$. The first term of this loss measures the distance (weighted by the precision Ω) between the linear predictor $\boldsymbol{\gamma}\tilde{X}$ and a future response $\tilde{Y}$. The second term promotes a sparse optimal summary, $\boldsymbol{\gamma}$. The penalty parameter λ determines the relative importance of these two components. Expanding the quadratic form gives:

$$\begin{aligned} \mathcal{L}_\lambda(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}) &= \frac{1}{2}\left(\tilde{Y}^T \Omega \tilde{Y} - 2\tilde{X}^T \boldsymbol{\gamma}^T \Omega \tilde{Y} + \tilde{X}^T \boldsymbol{\gamma}^T \Omega \boldsymbol{\gamma} \tilde{X} \right) + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_0 \\ &= \left(\tilde{X}^T \boldsymbol{\gamma}^T \Omega \boldsymbol{\gamma} \tilde{X} - 2\tilde{X}^T \boldsymbol{\gamma}^T \Omega \tilde{Y} \right) + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_0 + \text{constant}. \end{aligned} \tag{9}$$

Integrating over $(\tilde{Y}, \tilde{X}, \Theta \mid \mathbf{Y}, \mathbf{X})$ (and dropping the constant) gives:

$$\begin{aligned} \mathcal{L}_\lambda(\boldsymbol{\gamma}) &= \mathbb{E}[\mathcal{L}_\lambda(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma})] \\ &= \mathbb{E}\left[\mathrm{tr}[\boldsymbol{\gamma}^T \Omega \boldsymbol{\gamma} \tilde{X} \tilde{X}^T] \right] - 2\mathbb{E}\left[\mathrm{tr}[\boldsymbol{\gamma}^T \Omega \tilde{Y} \tilde{X}^T] \right] + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_0, \\ &= \mathbb{E}\left[\mathrm{tr}[\boldsymbol{\gamma}^T \Omega \boldsymbol{\gamma} S] \right] - 2\mathrm{tr}[A\boldsymbol{\gamma}^T] + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_0, \\ &= \mathrm{tr}[M\boldsymbol{\gamma} S \boldsymbol{\gamma}^T] - 2\mathrm{tr}[A\boldsymbol{\gamma}^T] + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_0, \end{aligned} \tag{10}$$

where

$$\begin{aligned} A &\equiv \mathbb{E}[\Omega \tilde{Y} \tilde{X}^T], \\ S &\equiv \mathbb{E}[\tilde{X}\tilde{X}^T] = \overline{\Sigma}_x, \\ M &\equiv \overline{\Omega}, \end{aligned} \tag{11}$$

and the overlines denote posterior means. Define the Cholesky decompositions $M = LL^T$ and $S = QQ^T$. To make the optimization problem tractable we replace the ℓ_0 norm with the ℓ_1 norm, leading to an expression that can be formulated in the form of a standard penalized regression problem:

$$\mathcal{L}_\lambda(\boldsymbol{\gamma}) = \left\| \left[Q^T \otimes L^T \right] \mathbf{vec}(\boldsymbol{\gamma}) - \mathbf{vec}(L^{-1} A Q^{-T}) \right\|_2^2 + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_1, \tag{12}$$

with covariates $Q^T \otimes L^T$, "data" $L^{-1}AQ^{-T}$, and regression coefficients $\boldsymbol{\gamma}$ (see the Appendix for details). Accordingly, (12) can be optimized using existing software such as the `lars` R package of Efron et al. (2004) and still yield sparse solutions.

2.2. Sparsity-utility trade-off plots

Rather than attempting to determine an "optimal" value of λ, we advocate displaying plots that reflect the utility attenuation due to λ-induced sparsification. We define the "loss gap" between a λ-sparse solution, $\mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}_\lambda^*)$, and the optimal unpenalized (non-sparse, $\lambda = 0$) summary, $\mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}^*)$ as

$$\Delta_\lambda = \mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}_\lambda^*) - \mathcal{L}(\tilde{Y}, \tilde{X}, \Theta, \boldsymbol{\gamma}^*). \tag{13}$$

As a function of $(\tilde{Y}, \tilde{X}, \Theta)$, Δ_λ is itself a random variable which we can sample by obtaining posterior draws from $p(\tilde{Y}, \tilde{X}, \Theta \mid \mathbf{Y}, \mathbf{X})$. The posterior distribution(s) of Δ_λ (for various λ) therefore reflects the deterioration in utility attributable to "sparsification". Plotting these distributions as a function of λ allows one to

visualize this trade-off. Specifically, $\pi_\lambda \equiv \Pr(\Delta_\lambda < 0 \mid \mathbf{Y}, \mathbf{X})$ is the (posterior) probability that the λ-sparse summary is no worse than the non-sparse summary. Using this framework, a useful heuristic for obtaining a single sparse summary is to report the sparsest model (associated with the highest λ) such that π_λ is higher than some pre-determined threshold, κ; we adopt this approach in our application section.

We propose summarizing the posterior distribution of Δ_λ via two types of plots. First, one can examine posterior means and credible intervals of Δ_λ for a sequence of models indexed by λ. Similarly, one can plot π_λ across the same sequence of models. Also, for a fixed value of λ, one can produce graphs where nodes represent predictor variables and response variables and an edge is drawn between nodes whenever the corresponding element of γ^*_λ is non-zero. All three types of plots are exhibited in Section 3.

2.3. Relation to previous methods

Loss function (12) is similar in form to the univariate *DSS* (decoupled shrinkage and selection) strategy developed by Hahn and Carvalho (2015). Our approach generalizes Hahn and Carvalho (2015) by optimizing over the matrix $\boldsymbol{\gamma} \in \mathbb{R}^{qxp}$ rather than a single vector of regression coefficients, extending the sparse summary utility approach to seemingly unrelated regression models (Brown, Vannucci and Fearn, 1998; Wang, 2010).

Additionally, the present method considers random predictors, $\tilde{X}$, whereas Hahn and Carvalho (2015) considered only a matrix of fixed design points. The impact of accounting for random predictors on the posterior summary variable selection procedure is examined in more detail in the application section.

An important difference between the sparse summary utility approach and previous approaches is in the role played by the posterior distribution. Many Bayesian variable importance metrics are based on the posterior distribution of an indicator variable that records if a given variable is non-zero (included in the model). The model we will use in our application utilizes such an indicator vector, which is called α. For example, a widely-used model selection heuristic is to examine the "inclusion probability" of predictor i, defined as the posterior mean of component α_i. However, any approach based on the posterior mean of α necessarily ignores information about the codependence between its elements, which can be substantial in cases of collinear predictors. Our method focuses instead on the expected log-density of future predictions, which synthesizes information from all parameters simultaneously in gauging how important they are in terms of future predictions.

3. Applications

In this section, the sparse posterior summary method is applied to a data set from the finance (asset pricing) literature. A key component of our analysis will be a comparison between the posterior summaries obtained when the predictors are drawn at random versus when they are assumed fixed.

The response variables are returns on 25 tradable portfolios and our predictor variables are returns on 10 other portfolios thought to be of theoretical importance. In the asset pricing literature Ross (1976), the response portfolios represent assets to be priced (so-called *test assets*) and the predictor portfolios represent distinct sources of variation (so-called *risk factors*). More specifically, the test assets Y represent trading strategies based on company size (total value of stock shares) and book-to-market (the ratio of the company's accounting valuation to its size); see Fama and French (1992) and Fama and French (2015) for details. Roughly, these

assets serve as a lower-dimensional proxy for the stock market. The risk factors are also portfolios, but ones which are thought to represent *distinct* sources of risk. What constitutes a distinct source of risk is widely debated, and many such factors have been proposed in the literature (Cochrane, 2011). We use monthly data from July 1963 through February 2015, obtained from Ken French's website:

`http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/`.

Our analysis investigates which subset of risk factors are most relevant (as defined by our utility function). As our initial candidates, we consider factors known in previous literature as: market, size, value, direct profitability, investment, momentum, short term reversal, long term reversal, betting against beta, and quality minus junk. Each factor is constructed by cross-sectionally sorting stocks by various characteristics of a company and forming linear combinations based on these sorts. For example, the value factor is constructed using the book-to-market ratio of a company. A high ratio indicates the company's stock is a "value stock" while a low ratio leads to a "growth stock" assessment. Essentially, the value factor is a portfolio built by going long stocks with high book-to-market ratio and shorting stocks with low book-to-market ratio. For detailed definitions of the first five factors, see Fama and French (2015). In the figures to follow, each is labeled as, for example, "Size2 BM3," to denote the portfolio buying stocks in the second quintile of size and the third quintile of book-to-market ratio.

Recent related work includes Ericsson and Karlsson (2004) and Harvey and Liu (2015). Ericsson and Karlsson (2004) follow a Bayesian model selection approach based off of inclusion probabilities, representing the preliminary inference step of our methodology. Harvey and Liu (2015) take a different approach that utilizes multiple hypothesis testing and bootstrapping.

3.1. Results

As described in Section 1.1, the first step of our analysis consists of fitting a Bayesian model. We fit model (1) using a variation of the well-known stochastic search variable selection algorithm of George and McCulloch (1993) and similar to Brown, Vannucci and Fearn (1998) and Wang (2010). Details are given in the Appendix.

In the subsections to follow, we will show the following two figures. First, we plot the expectation of Δ_λ (and associated posterior credible interval) across a range of λ penalties. Recall, Δ_λ is the "loss gap" between a sparse summary and the best non-sparse (saturated) summary, meaning that smaller values are "better". Additionally, we plot the probability that a given model is no worse than the saturated model π_λ on this same figure, where "no worse" means $\Delta_\lambda < 0$. Note that even for very weak penalties (small λ), the distribution of Δ_λ will have non-zero variance and therefore even if it is centered about zero, some mass can be expected to fall above zero; practically, this means that $\pi_\lambda > 0.5$ is a very high score.

Second, we display a summary graph of the selected variables for the $\kappa = 12.5\%$ threshold. Recall that this is the highest penalty (sparsest graph) that is no worse than the saturated model with 12.5% posterior probability. For these graphs, the response and predictor variables are colored gray and white, respectively. A test asset label of, for example, "Size2 BM3," denotes the portfolio that buys stocks in the second quintile of size and the third quintile of book-to-market ratio. The predictors without connections to the responses under the optimal graph are not displayed.

These two figures are shown in four scenarios:

1. Random predictors.
2. Fixed predictors.
3. Random predictors under alternative prior.
4. Fixed predictors under alternative prior.

The "alternative prior" scenario serves to show the impact of the statistical modeling comprising *step 1*. Specifically, we use the same Monte Carlo model fitting procedure as before (described in the Appendix) but fix α to the identity vector. That is, we omit the point-mass component of the priors for the elements of $\boldsymbol{\beta}$.

3.1.1. Random predictors

This section introduces our baseline example where the risk factors (predictors) are random. We evaluate the set of potential models by analyzing plots such as figure 1. This shows Δ_λ and π_λ evaluated across a range of λ values. Additionally, we display the posterior uncertainty in the Δ_λ metric with gray vertical uncertainty bands: these are the centered $P\%$ posterior credible intervals where $\kappa = (1 - P)/2$. As the accuracy of the sparsified solution increases, the posterior of Δ_λ concentrates around zero by construction, and the probability of the model being no worse than the saturated model, π_λ, increases. We choose the sparsest model such that its corresponding $\pi_\lambda > \kappa = 12.5\%$. This model is displayed in figure 2 and is identified by the black dot in figure 1.

The selected set of factors in graph 2 are the market (Mkt.RF), value (HML), and size (SMB). This three factor model is no worse than the saturated model with 12.5% posterior probability where all test assets are connected to all risk factors. Note also that in our selected model almost every test asset is distinctly tied to one of either value or size and the market factor. These are the three factors of Ken French and Eugene Fama's pricing model developed in Fama and French (1992). They are known throughout the finance community as being "fundamental dimensions" of the financial market, and our procedure is consistent with this widely held belief at a small κ level.

The characteristics of the test assets in graph 2 are also important to highlight. The test portfolios that invest in small companies ("Size1" and "Size2") are primarily connected to the SMB factor which is designed as a proxy for the risk of small companies. Similarly, the test portfolios that invest in high book-to-market companies ("BM4" and "BM5") have connections to the HML factor which is built on the idea that companies whose book value exceeds the market's perceived value should generate a distinct source of risk. As previously noted, all of the test portfolios are connected to the market factor suggesting that it is a relevant predictor even for the sparse $\kappa = 12.5\%$ selection criterion.

In figure 3, we examine how different choices of the κ threshold change the selected set of risk factors. In this analysis, there is a tradeoff between the posterior probability of being "close" to the saturated model and the utility's preference for sparsity. When the threshold is low ($\kappa = 2, 4$, and 12.5%) the summarization procedure selects relatively sparse graphs with up to three factors (Mkt.RF, HML, and SMB). The market (Mkt.RF) and size (SMB) factors appear first, connected to a small number of the test assets ($\kappa = 2\%$). As the threshold is increased, the point summary becomes denser and correspondingly more predictively accurate (as measured by the utility function). The value factor (HML) enters at $\kappa = 12.5\%$ and quality minus junk (QMJ), investment (CMA), and profitability (RMW) factors enter at $\kappa = 32.5\%$. The graph for $\kappa = 32.5\%$ excluding QMJ is essentially the new five factor model proposed by Fama and French (2015). The five Fama-French factors (plus OMJ with three connections) persist up to the $\kappa = 47.5\%$ threshold.

This indicates that, up to a high posterior probability, the five factor model of Fama and French (2015) does no worse than an asset pricing model with all ten factors connected to all test assets.

Notice also that our summarization procedure displays the specific relationship between the factors and test assets through the connections. Using this approach, the analyst is able to identify which predictors drive variation in which responses and at what thresholds they may be relevant. This feature is significant for summarization problems where individual characteristics of the test portfolios and their joint dependence on the risk factors is may be a priori unclear.

As κ approaches the 50% threshold ($\kappa = 49.75\%$ in figure 3), the model summary includes all ten factors. Requesting a summary with this level of certainty results in little sparsification. However, compared to the nearby $\kappa = 47.5\%$ model with only six factors, we also now know that the remaining four contribute little to our utility. These factors are betting against beta (BAB), momentum (Mom), long term reversal (LTR), and short term reversal (STR). Sparse posterior summarization applied in this context allows an analyst to study the impact of risk factors on pricing while taking uncertainty into account. Coming to a similar conclusion via common alternative techniques (e.g., component-wise ordinary least squares combined with thresholding by t-statistics) is comparatively ad hoc; our method is simply a perspicuous summary of a posterior distribution. Likewise, applying sparse regression techniques based on ℓ_1 penalized likelihood methods would not take into account the residual correlation Ψ, nor would that approach naturally accommodate random predictors.

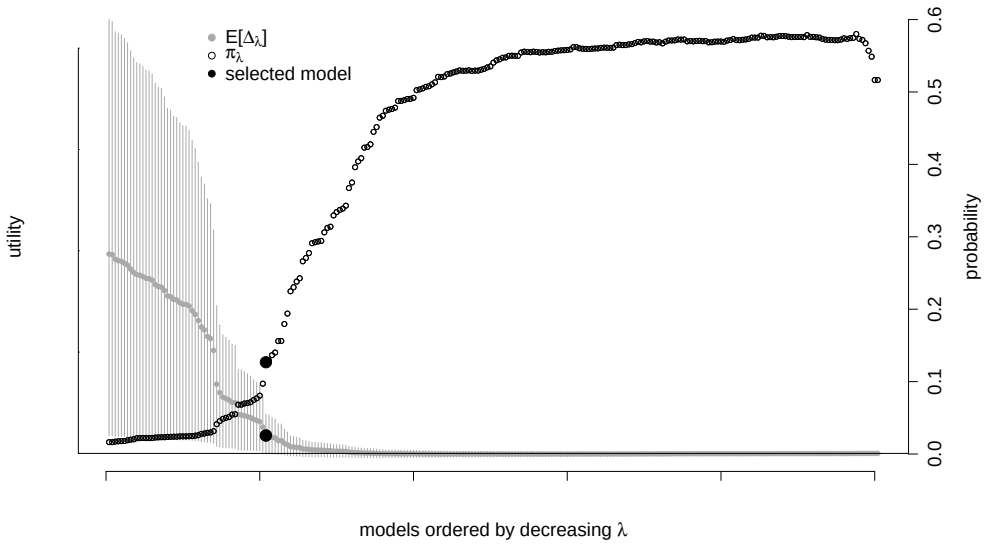

FIG 1. *Evaluation of Δ_λ and π_λ along the solution path for the 25 size/value portfolios modeled by the 10 factors. An analyst may use this plot to select a particular model. Uncertainty bands are 75% posterior intervals on the Δ_λ metric. The large black dot represents the model selected in Figure 2.*

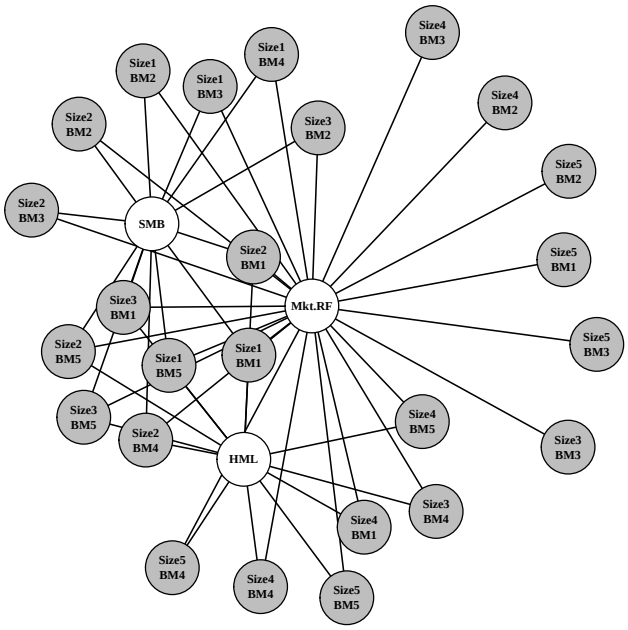

FIG 2. *The selected model for 25 size/value portfolios modeled by the 10 factors. The responses and predictors are colored in gray and white, respectively. Edges represent nonzero components of the optimal action, $\boldsymbol{\gamma}$.*

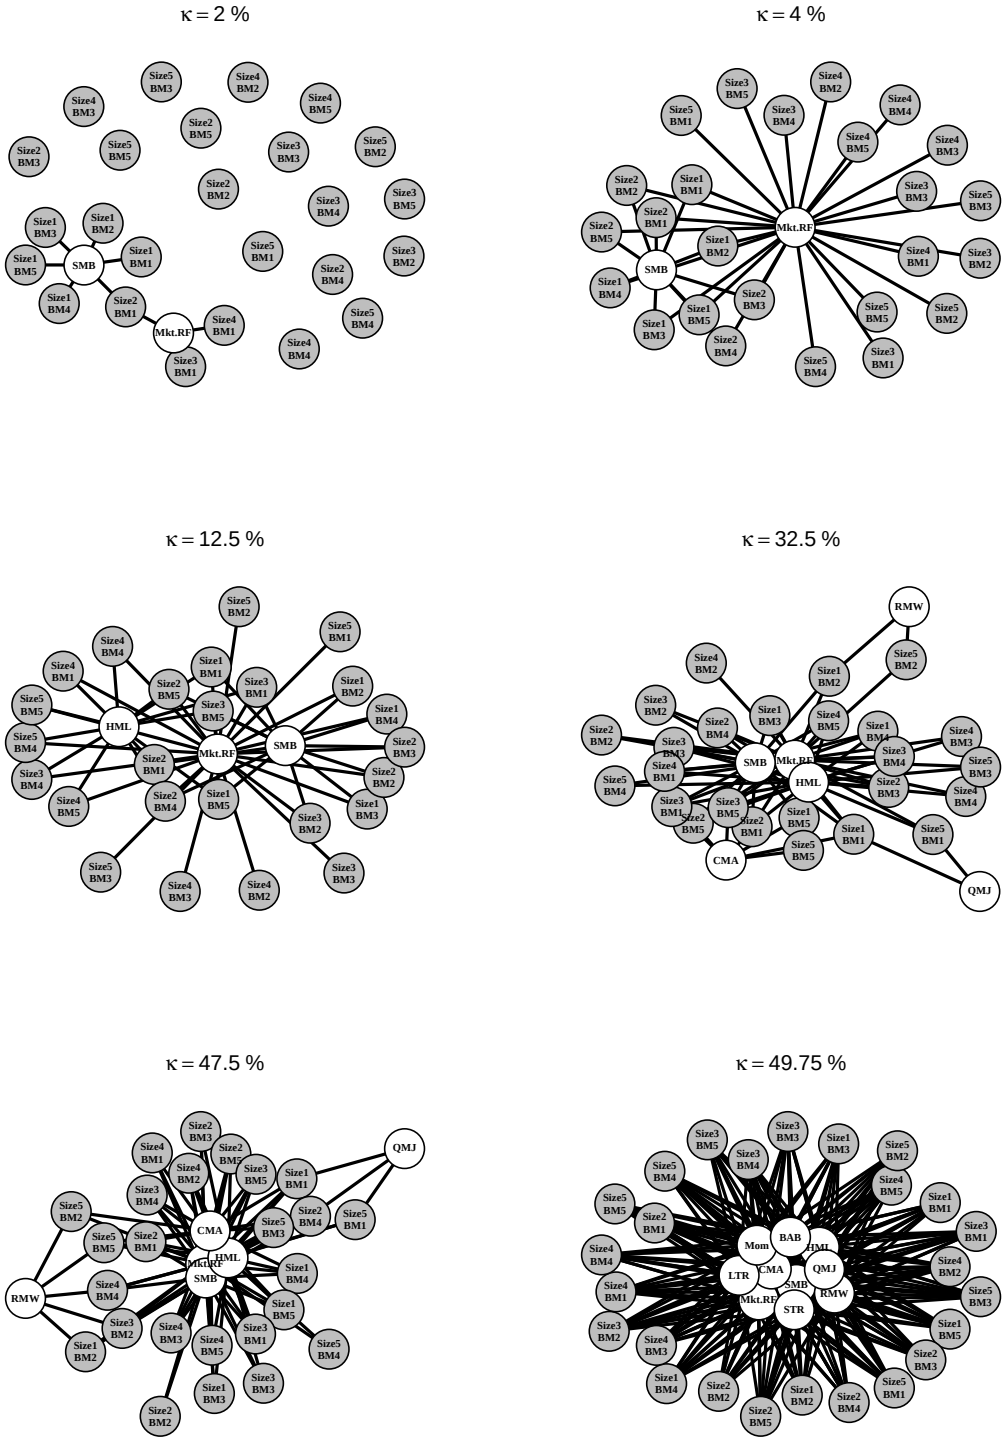

FIG 3. *Sequence of selected models for varying threshold level* κ *under the assumption of* ***random predictors****.*

3.1.2. Fixed predictors

In this section, we consider posterior summarization with the loss function derived under the assumption of *fixed predictors*. The analogous loss function when the predictor matrix is fixed is:

$$\mathcal{L}_\lambda(\boldsymbol{\gamma}) = \left\| \left[Q_f^T \otimes L^T \right] \mathbf{vec}(\boldsymbol{\gamma}) - \mathbf{vec}(L^{-1} A_f Q_f^{-T}) \right\|_2^2 + \lambda \left\| \mathbf{vec}(\boldsymbol{\gamma}) \right\|_1, \tag{14}$$

with $Q_f Q_f^T = \mathbf{X}^T\mathbf{X}$, $A_f = \mathbb{E}[\Omega \tilde{\mathbf{Y}}^T \mathbf{X}]$, and $M = \overline{\Omega} = LL^T$; compare to (11) and (12). The derivation of (14) is similar to the presentation in Section 2 and may be found in the Appendix. The corresponding version of the loss gap is

$$\Delta_\lambda = \mathcal{L}(\tilde{\mathbf{Y}}, \mathbf{X}, \Theta, \boldsymbol{\gamma}_\lambda^*) - \mathcal{L}(\tilde{\mathbf{Y}}, \mathbf{X}, \Theta, \boldsymbol{\gamma}^*). \tag{15}$$

which has distribution induced by the posterior over $(\tilde{\mathbf{Y}}, \Theta)$ rather than $(\tilde{Y}, \tilde{X}, \Theta)$ as before. By fixing $\mathbf{X}$, the posterior of Δ_λ has smaller dispersion which results in denser summaries for the same level of κ. For example, compare how dense Figure 5 is relative to Figure 2. The denser graph in Figure 5 contains nine out of ten potential risk factors compared to just three in Figure 2, which correspond to the Fama-French factors described in Fama and French (1992). Recall, both graphs represent the sparsest model such that the probability of being no worse than the saturated model is greater than $\kappa = 12.5\%$ — the difference is that one of the graphs defines "worse-than" in terms of a fixed set of risk factor returns while the other acknowledge that those returns are themselves uncertain in future periods.

Figure 6 demonstrates this problem for several choices of the uncertainty level. Regardless of the uncertainty level chosen, the selected models contain most of the ten factors and many edges. In fact, it is difficult to distinguish even the $\kappa = 2\%$ and $\kappa = 49.75\%$ models.

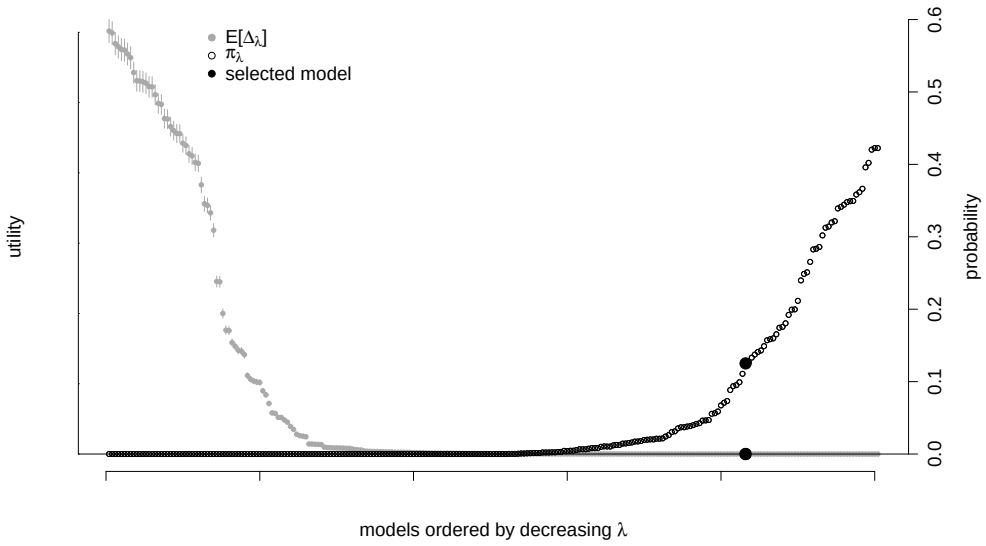

FIG 4. *Evaluation of Δ_λ and π_λ along the solution path for the 25 size/value portfolios modeled by the 10 factors assuming* ***fixed predictors (factors)****. An analyst may use this plot to select a particular model. Uncertainty bands are 75% posterior intervals on the Δ_λ metric. The large black dot represents the model selected in Figure 5.*

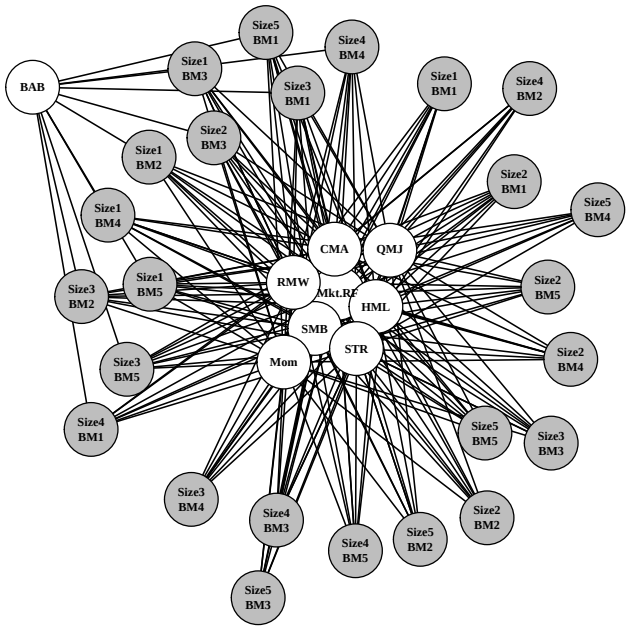

FIG 5. *The selected model for 25 size/value portfolios modeled by the 10 factors when* ***uncertainty in future factor returns is not taken into account****. The responses and predictors are colored in gray and white, respectively. Edges represent nonzero components of the optimal action, $\boldsymbol{\gamma}$.*

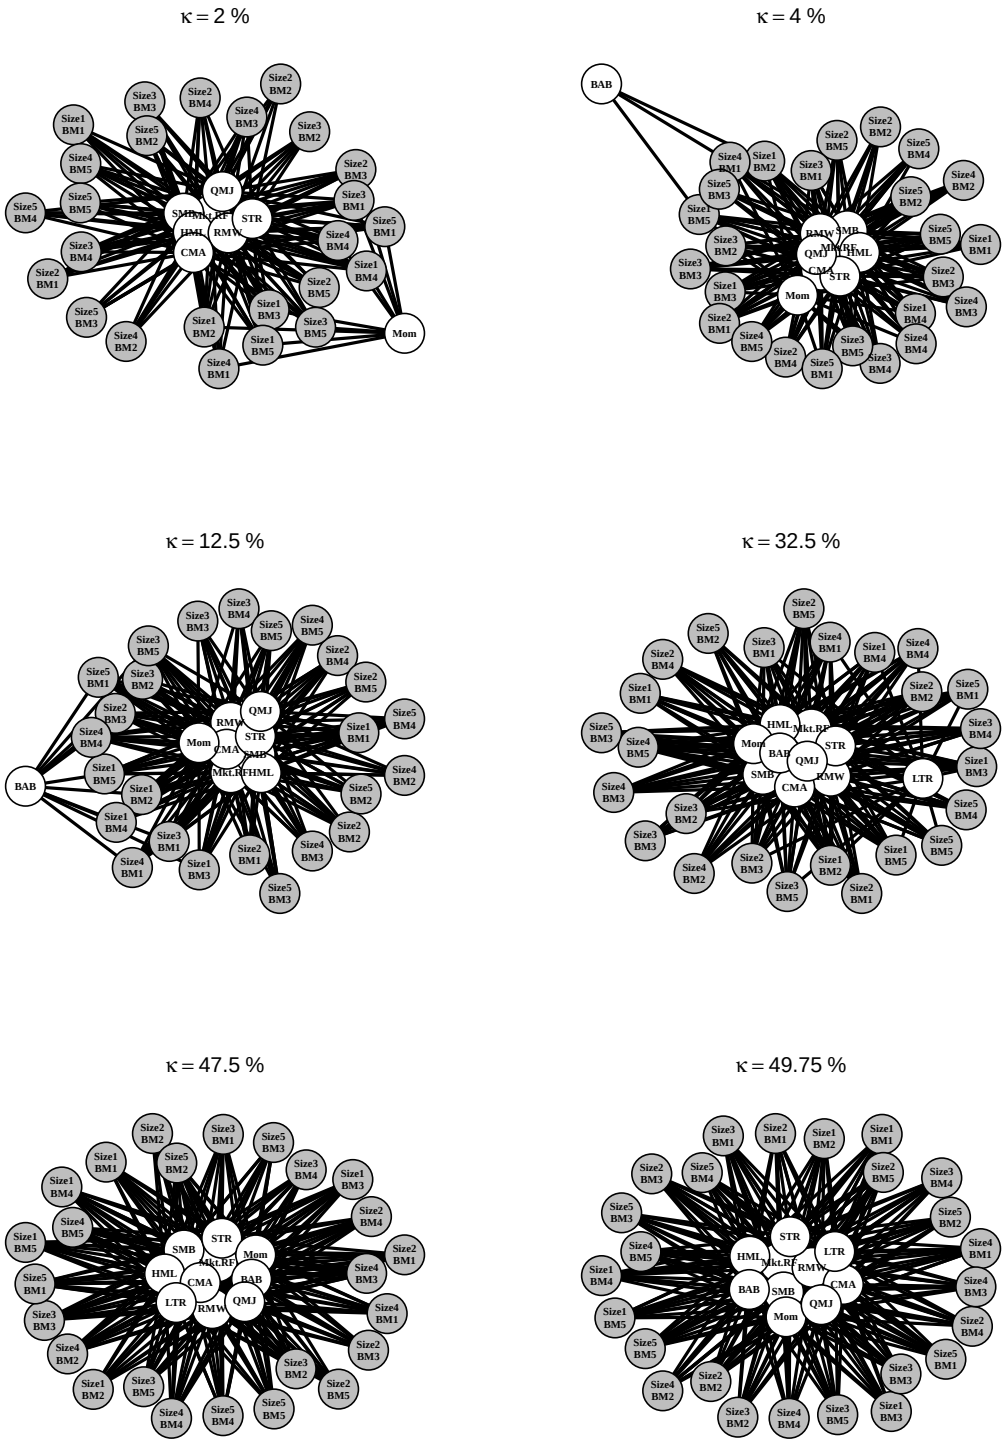

FIG 6. *Sequence of selected models for varying threshold level κ under the assumption of* ***fixed predictors****.*

3.1.3. Alternative prior analysis

Here, we consider how our posterior summaries change as a function of using a different posterior, based on a different choice of prior. Specifically, in this section we do not employ model selection point-mass priors on the elements of $\boldsymbol{\beta}$ as we did in the above analysis. These results are displayed in figures 7 and 8. Broadly, the same risk factors are flagged as important — the market factor followed by the size (SMB) and value (HML) factors. One notable difference is that the quality minus junk (QMJ), investment (CMA), and profitability (RMW) factors appear at smaller levels of κ. This result is intuitive in the sense that point-mass priors demand stronger evidence for a variable to impact the posterior means defining the loss function. Without the strong shrinkage imposed by the point-mass priors, these risk factors show up more strongly in the posterior and hence in the posterior summary. In each case, the three Fama and French factors from Fama and French (1992) predictably appear and seem to be the only relevant factors for pricing these 25 portfolios.

Similarly, the weaker shrinkage model in the fixed predictor version (figures 9 and 10) yields yet denser summaries (for a given level of κ).

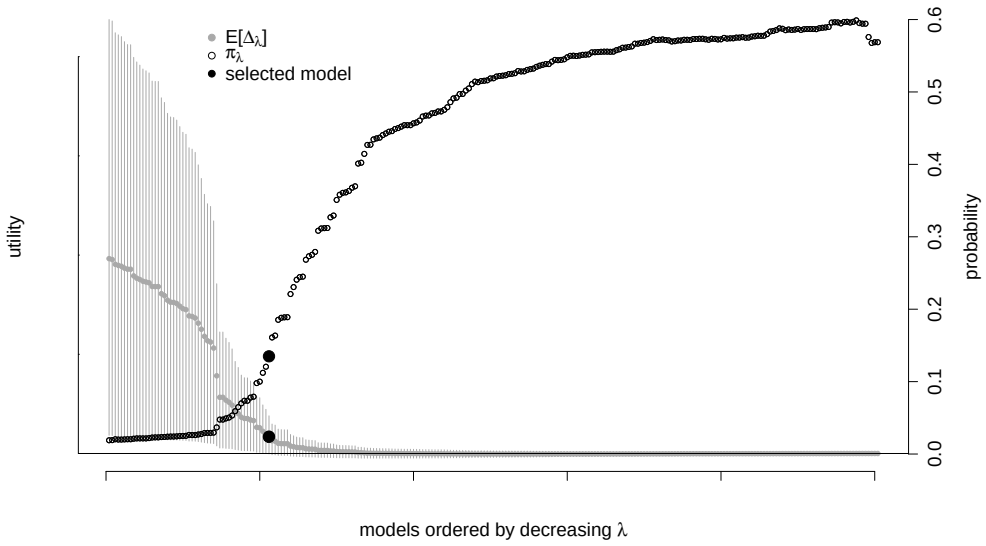

FIG 7. *Evaluation of Δ_λ and π_λ along the solution path for the 25 size/value portfolios modeled by the 10 factors with alternative prior. An analyst may use this plot to select a particular model. Uncertainty bands are 75% posterior intervals on the Δ_λ metric. The large black dot represents the model selected in Figure 8.*

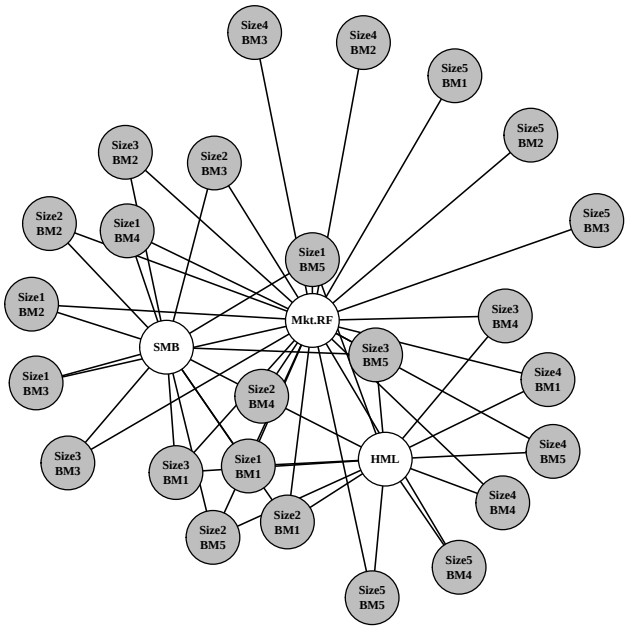

FIG 8. *The selected model for 25 size/value portfolios modeled by the 10 factors with alternative prior. The responses and predictors are colored in gray and white, respectively. Edges represent nonzero components of the optimal action, $\boldsymbol{\gamma}$.*

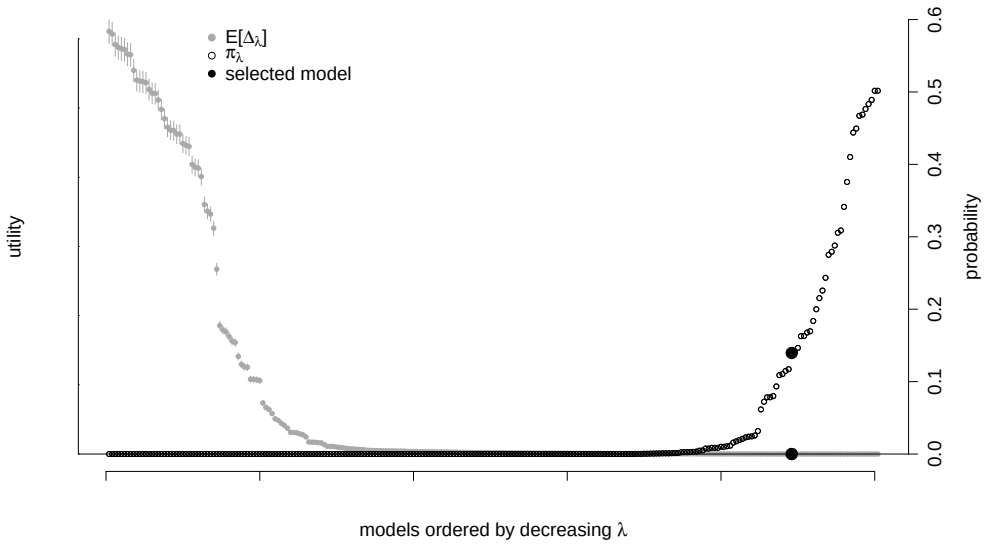

FIG 9. *Evaluation of Δ_λ and π_λ along the solution path for the 25 size/value portfolios modeled by the 10 factors assuming* ***fixed predictors (factors)*** *and with alternative prior. An analyst may use this plot to select a particular model. Uncertainty bands are 75% posterior intervals on the Δ_λ metric. The large black dot represents the model selected in Figure 10.*

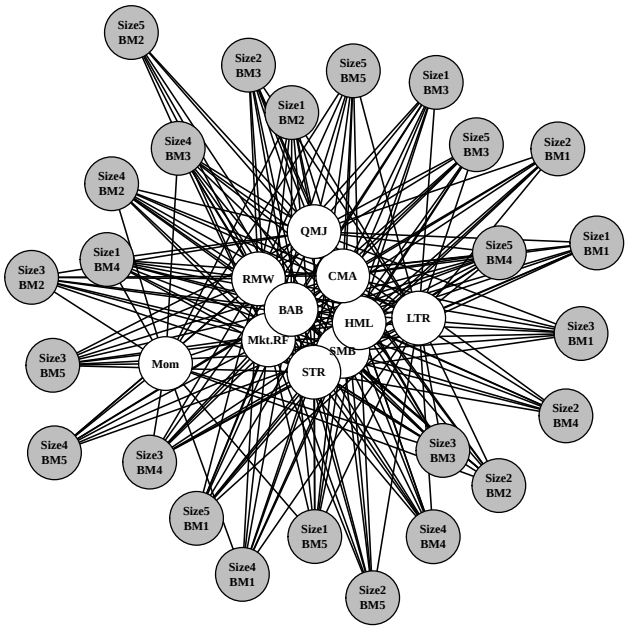

FIG 10. *The selected model for 25 size/value portfolios modeled by the 10 factors when* ***uncertainty in future factor returns is not taken into account*** *and with alternative prior. The responses and predictors are colored in gray and white, respectively. Edges represent nonzero components of the optimal action, $\boldsymbol{\gamma}$.*

3.1.4. Comparison of four scenarios at fixed κ

The selected summary graphs for the four scenarios are displayed together for comparison in figure 11. Observe that graphs (c) and (d) selected under the alternative prior are marginally denser than their counterparts (a) and (b) under the point-mass model selection prior. However, the assumption of random predictors results in notably sparser summaries – graphs (a) and (c) are much sparser than (b) and (d). These comparisons emphasize the impact that incorporating random predictors may have on a variable selection procedure; especially the present approach where we extract point summaries from a posterior by utilizing uncertainty in all unknowns $(\tilde{Y}, \tilde{X}, \Theta)$.

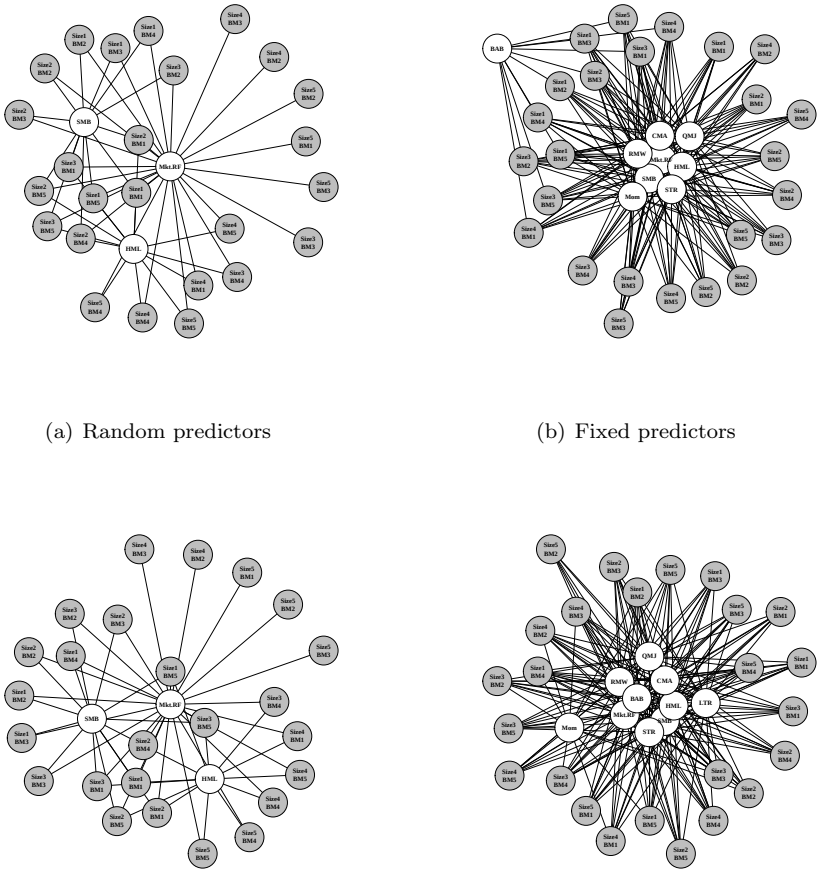

(a) Random predictors

(b) Fixed predictors

(c) Random predictors and alternative prior

(d) Fixed predictors and alternative prior

FIG 11. *Comparison of selected models under four scenarios.*

4. Conclusion

In this paper, we propose a general model selection procedure for multivariate linear models when future realizations of the predictors are unknown. Such models are widely used in many areas of science and economics, including genetics and asset

pricing. Our utility-based sparse posterior summary procedure is a multivariate extension of the "decoupling shrinkage and selection" methodology of Hahn and Carvalho (2015). The approach we develop has three steps: (*i*) fit a Bayesian model, (*ii*) specify a utility function with a sparsity-inducing penalty term and optimize its expectation, and (*iii*) graphically summarize the posterior impact (in terms of utility) of the sparsity penalty. Our utility function is based on the kernel of the conditional distribution responses given the predictors and can be formulated as a tractable convex program. We demonstrate how our procedure may be used in asset pricing under a variety of modeling choices.

The remainder of this discussion takes a step back from the specifics of the seemingly unrelated regressions model and considers a broader role for utility-based posterior summaries.

A paradox of applied Bayesian analysis is that posterior distributions based on relatively intuitive models like the SUR model are often just as complicated as the data itself. For Bayesian analysis to become a routine tool for practical inquiry, methods for summarizing posterior distributions must be developed apace with the models themselves. A natural starting point for developing such methods is decision theory, which suggests developing loss functions specifically geared towards practical posterior summary. As a matter of practical data analysis, articulating an apt loss function has been sorely neglected relative to the effort typically lavished on the model specification stage, specifically prior specification. Ironically (but not surprisingly) our application demonstrates that one's utility function has a dominant effect on the posterior summaries obtained relative to which prior distribution is used.

This paper makes two contributions to this area of "utility design". First, we propose that the likelihood function has a role to play in posterior summary apart from its role in inference. That is, one of the great practical virtues of likelihood-based statistics is that the likelihood serves to summarize the data by way of the corresponding point estimates. By using the log-density as our utility function applied to *future* data, we revive the fundamental summarizing role of the likelihood. Additionally, note that this approach allows three distinct roles for parameters. First, all parameters of the model appear in defining the posterior predictive distribution. Second, some parameters appear in *defining* the loss function; Ψ plays this role in our analysis. Third, some parameters define the action space. In this framework there are no "nuisance" parameters that vanish from the estimator as soon as a marginal posterior is obtained. Once the likelihood-based utility is specified, it is a natural next step to consider augmenting the utility to enforce particular features of the desired point summary. For example, our analysis above was based on a utility that explicitly rewards sparsity of the resulting summary. A traditional instance of this idea is the definition of high posterior density regions, which are defined as the *shortest, contiguous* interval that contains a prescribed fraction of the posterior mass.

Our second contribution is to consider not just one, but a range, of utility functions and to examine the posterior distributions of the corresponding posterior loss. Specifically, we compare the utility of a sparsified summary to the utility of the optimal non-sparse summary. Interestingly, these utilities are random variables themselves (defined by the posterior distribution) and examining their distributions provides a fundamentally Bayesian way to measure the extent to which the sparsity preference is driving one's conclusions. The idea of comparing a hypothetical continuum of decision-makers based on the posterior distribution of their respective utilities represents a principled Bayesian approach to exploratory data analysis. This is an area of ongoing research.

Appendix A: Matrix-variate Stochastic Search

A.1. Model fitting: The marginal and conditional distributions

The future values of the response and covariates are unknown. Acknowledging this uncertainty is important in the overall decision of which covariates to select and is a necessary ingredient of the selection procedure. As our examples consider financial asset return data, we choose to model the marginal distribution of the covariates via a latent factor model detailed in Murray et al. (2013). The responses are modeled conditionally on the covariates via a matrix-variate stochastic search which is a multivariate of extension of stochastic search variable selection (SSVS) from George and McCulloch (1993). Recalling the block structure for the covariance of the full joint distribution of (X, Y):

$$\Sigma = \left[\begin{array}{c|c} \boldsymbol{\beta}^T\Sigma_x\boldsymbol{\beta} + \Psi & (\Sigma_x\boldsymbol{\beta})^T \\ \hline \Sigma_x\boldsymbol{\beta} & \Sigma_x \end{array}\right], \tag{16}$$

we obtain posterior samples of Σ by sampling the conditional model parameters using a matrix-variate stochastic search algorithm (described below) and sampling the covariance of X from a latent factor model where it is marginally normally distributed. To reiterate our procedure is

- Σ_x is sampled from independent latent factor model,
- $\boldsymbol{\beta}$ is sampled from matrix-variate MCMC,
- Ψ is sampled from matrix-variate MCMC.

A.1.1. Modeling a full residual covariance matrix

In order to sample a full residual covariance matrix, we augment the predictor matrix with a latent factor f by substituting $\epsilon_j = b_j f + \tilde{\epsilon}_j$:

$$Y_j = \beta_{j1}X_1 + \cdots + \beta_{jp}X_p + b_j f + \tilde{\epsilon}_j, \qquad \tilde{\boldsymbol{\epsilon}} \sim \mathrm{N}(0, \tilde{\Psi}), \tag{17}$$

where $\tilde{\Psi}$ is now diagonal. Assuming that $f \sim N(0, 1)$ is shared among all response variables j and $\mathbf{b} \in \mathbb{R}^{qx1}$ is a vector of all coefficients b_j, the total residual variance may be expressed as:

$$\Psi = \mathbf{b}\mathbf{b}^T + \tilde{\Psi}. \tag{18}$$

We incorporate this latent factor model into the matrix-variate MCMC via a simple Gibbs step to draw posterior samples of f. This augmentation allows us to draw samples of Ψ that are not constrained to be diagonal.

A.1.2. Modeling the marginal distribution: A latent factor model

We model covariates via a latent factor model of the form:

$$\begin{gathered} \mathbf{X}_t = \mu_x + \mathbf{B}\mathbf{f}_t + \mathbf{v}_t \\ \mathbf{v}_t \sim \mathrm{N}(0, \boldsymbol{\Lambda}), \quad \mathbf{f}_t \sim \mathrm{N}(0, \mathbb{I}_k), \quad \mu_x \sim \mathrm{N}(0, \Phi) \end{gathered} \tag{19}$$

where Λ is assumed diagonal and the set of k latent factors f_t are independent. The covariance of the covariates is constrained by the factor decomposition and takes the form:

$$\Sigma_x = \mathbf{B}\mathbf{B}^T + \Lambda. \tag{20}$$

Recall that this is only a potential choice for the $p(X)$ and it is chosen here primarily motivated by the applied context where financial assets tend to depend across each other through common factors. Our variable selection procedure would follow if any other choice was made at this point. To estimate this model, a convenient, efficient choice is the R package `bfa` (Murray, 2015). The software allows us to sample the marginal covariance as well as the marginal mean via a simple Gibbs step assuming a normal prior on μ_x.

A.1.3. Modeling the conditional distribution: A matrix-variate stochastic search

We model the conditional distribution, $Y|X$, by developing a multivariate extension of stochastic search variable selection of George and McCulloch (1993). Recall that the conditional model is: $\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} \sim \mathcal{N}(\mathbb{I}_{N\times N},\ \Psi_{q\times q})$. In order to sample different subsets of covariates (different models) during the posterior simulations, we introduce an additional parameter $\alpha \in \mathbb{R}^p$ that is a binary vector identifying a particular model. In other words, all entries i for which $\alpha_i = 1$ denote covariate i as included in model M_α. Specifically, we write the model identified by α as $M_\alpha : \mathbf{Y} - \mathbf{X}_\alpha\boldsymbol{\beta}_\alpha \sim \mathcal{N}(\mathbb{I}_{N\times N},\ \Psi_{q\times q})$. As in George and McCulloch (1993), we aim to explore the posterior on the model space, $\mathbf{P}(M_\alpha \mid \mathbf{Y})$. Our algorithm explores this model space by calculating a Bayes factor for a particular model M_α. Given that the response $\mathbf{Y}$ is matrix instead of a vector, we derive the Bayes factor as a product of vector response Bayes factors. This is done by separating the marginal likelihood of the response matrix as a product of marginal likelihoods across the separate vector responses. This derivation requires our priors to be independent across the responses and is shown in the Appendix. It is important to note that we do not run a standard SSVS on each univariate response regression separately. Instead, we generalize George and McCulloch (1993) and require all covariates to be included or excluded from a model for each of the responses *simultaneously*.

The marginal likelihood requires priors for the parameters $\boldsymbol{\beta}$ and σ parameters in our model. We choose the standard g-prior for linear models because it permits an analytical solution for the marginal likelihood integral (Zellner, 1986; Zellner and Siow, 1984; Liang et al., 2008a).

Our Gibbs sampling algorithm directly follows the stochastic search variable selection procedure described in George and McCulloch (1993) using these calculated Bayes factors, now adapted to a multivariate setting. The aim is to scan through all possible covariates and determine which ones to include in the model identified through the binary vector α. At each substep of the MCMC, we consider an individual covariate i within a specific model and compute its inclusion probability as a function of the model's prior probability and the Bayes factors:

$$p_i = \frac{B_{a0}\mathbf{P}(M_{\alpha_a})}{B_{a0}\mathbf{P}(M_{\alpha_a}) + B_{b0}\mathbf{P}(M_{\alpha_b})}.$$

The Bayes factor B_{a0} is a ratio of marginal likelihoods for the model with covariate i included and the null model, and B_{b0} is the analogous Bayes factor for the model without covariate i. The prior on the model space, $\mathbf{P}(M_\gamma)$, can either be chosen to adjust for multiplicity or to be uniform - our results appear robust to both specifications. In this setting, adjusting for multiplicity amounts to putting equal prior mass on different sizes of models. In contrast, the uniform prior for models involving p covariates puts higher probability mass on larger models, reaching a

maximum for models with $\binom{p}{2}$ covariates included. The details of the priors on the model space and parameters, including an empirical Bayes choice of the g-prior hyperparameter, are discussed in the Appendix.

A.2. Details

Assume we have observed N realizations of data $(\mathbf{Y}, \mathbf{X})$. For model comparison, we calculate the Bayes factor with respect to the null model without any covariates. First, we calculate a marginal likelihood. This likelihood is obtained by integrating the full model over $\boldsymbol{\beta}_\alpha$ and σ multiplied by a prior, $\pi_\alpha(\boldsymbol{\beta}_\alpha, \sigma)$, for these parameters. A Bayes factor of a given model α versus the null model, $B_{\alpha 0} = \frac{m_\alpha(\mathbf{R})}{m_0(\mathbf{R})}$ with:

$$m_\alpha(\mathbf{Y}) = \int \text{Matrix Normal}_{N,q}\left(\mathbf{Y} \mid \mathbf{X}_\alpha \boldsymbol{\beta}_\alpha,\ \mathbb{I}_{NxN},\ \tilde{\Psi}_{q\times q}\right) \pi_\alpha(\boldsymbol{\beta}_\alpha, \sigma_i)\, d\boldsymbol{\beta}_\alpha d\sigma_i. \tag{21}$$

We assume independence of the priors across columns of $\mathbf{Y}$ so we can write the integrand in (21) as a product across each individual response vector:

$$\begin{aligned}
m_\alpha(\mathbf{Y}) &= \int \Pi_{i=1}^q\ N_N\left(\mathbf{Y}^i \mid \mathbf{X}_\alpha \boldsymbol{\beta}_\alpha^i,\ \sigma_i^2 \mathbb{I}_{NxN}\right) \pi_\alpha^i\left(\boldsymbol{\beta}_\alpha^i, \sigma_i\right) d\boldsymbol{\beta}_\alpha^i d\sigma_i \\
&\iff \\
m_\alpha(\mathbf{Y}) &= \int\ N_N\left(\mathbf{Y}^1 \mid \mathbf{X}_\alpha \boldsymbol{\beta}_\alpha^1,\ \sigma_1^2 \mathbb{I}_{NxN}\right) \pi_\alpha^1\left(\boldsymbol{\beta}_\alpha^1, \sigma_1\right) d\boldsymbol{\beta}_\alpha^1 d\sigma_1 \\
&\quad \times \cdots \times \int N_N\left(\mathbf{Y}^q \mid \mathbf{X}_\alpha \boldsymbol{\beta}_\alpha^q,\ \sigma_q^2 \mathbb{I}_{NxN}\right) \pi_\alpha^q\left(\boldsymbol{\beta}_\alpha^q, \sigma_q\right) d\boldsymbol{\beta}_\alpha^q d\sigma_q \\
&= m_\alpha\left(\mathbf{Y}^1\right) \times \cdots \times m_\alpha\left(\mathbf{Y}^q\right) \\
&= \Pi_{i=1}^q m_\alpha\left(\mathbf{Y}^i\right),
\end{aligned}$$

with:

$$\mathbf{Y}^i \sim N_N\left(\mathbf{X}_\alpha \boldsymbol{\beta}_\alpha^i,\ \sigma_i^2 \mathbb{I}_{NxN}\right). \tag{22}$$

Therefore, the Bayes factor for this matrix-variate model is just a product of Bayes factors for the individual multivariate normal models.

$$B_{\alpha 0} = \widetilde{B}_{\alpha 0}^1 \times \cdots \times \widetilde{B}_{\alpha 0}^q \tag{23}$$

with:

$$\widetilde{B}_{\alpha 0}^i = \frac{m_\alpha\left(\mathbf{Y}^i\right)}{m_0\left(\mathbf{Y}^i\right)}. \tag{24}$$

The simplification of the marginal likelihood calculation is crucial for analytical simplicity and for the resulting SSVS algorithm to rely on techniques already developed for univariate response models. In order to calculate the integral for each Bayes factor, we need priors on the parameters $\boldsymbol{\beta}_\alpha$ and σ. Since the priors are independent across the columns of $\mathbf{Y}$, we aim to define $\pi_\alpha^i\left(\boldsymbol{\beta}_\alpha^i, \sigma_i\right)\ \forall i \in \{1, ..., q\}$, which we express as the product: $\pi_\alpha^i(\sigma_i)\, \pi_\alpha^i\left(\boldsymbol{\beta}_\alpha^i \mid \sigma_i\right)$. Motivated by the work on regression problems of Zellner, Jeffreys, and Siow, we choose a non-informative prior for

σ_i and the popular g-prior for the conditional prior on $\boldsymbol{\beta}_\alpha^i$, (Zellner, 1986), (Zellner and Siow, 1980), (Zellner and Siow, 1984), (Jeffreys, 1961):

$$\pi_\alpha^i\left(\boldsymbol{\beta}_\alpha^i, \sigma_i \mid g\right) = \sigma_i^{-1}\mathrm{N}_{k_\alpha}\left(\boldsymbol{\beta}_\alpha^i \mid \mathbf{0}, g_\alpha^i\sigma_i^2(\mathbf{X}_\alpha^T(\mathbb{I} - N^{-1}\mathbf{1}\mathbf{1}^T)\mathbf{X}_\alpha)^{-1}\right). \tag{25}$$

Under this prior, we have an analytical form for the Bayes factor:

$$B_{\alpha 0} = \widetilde{B}_{\alpha 0}^1 \times \cdots \times \widetilde{B}_{\alpha 0}^q \tag{26}$$

$$= \Pi_{i=1}^q \frac{\left(1+g_\alpha^i\right)^{(N-k_\alpha-1)/2}}{\left(1+g_\alpha^i\frac{SSE_\alpha^i}{SSE_0^i}\right)^{(N+1)/2}}, \tag{27}$$

where SSE_α^i and SSE_0^i are the sum of squared errors from the linear regression of column $\mathbf{Y}^i$ on covariates $\mathbf{X}_\alpha$ and k_α is the number of covariates in model M_α. We allow the hyper parameter g to vary across columns of $\mathbf{Y}$ and depend on the model, denoted by writing, g_α^i.

We aim to explore the posterior of the model space, given our data:

$$\mathbf{P}\left(M_\alpha \mid \mathbf{Y}\right) = \frac{B_{\alpha 0}\mathbf{P}\left(M_\alpha\right)}{\Sigma_\alpha B_{\alpha 0}\mathbf{P}\left(M_\alpha\right)}, \tag{28}$$

where the denominator is a normalization factor. In the spirit of traditional stochastic search variable selection Garcia-Donato and Martinez-Beneito (2013), we propose the following Gibbs sampler to sample this posterior.

A.3. Gibbs Sampling Algorithm

Once the parameters $\boldsymbol{\beta}_\alpha$ and σ are integrated out, we know the form of the full conditional distributions for $\alpha_i \mid \alpha_1, \cdots, \alpha_{i-1}, \alpha_{i+1}, \cdots, \alpha_p$. We sample from these distributions as follows:

1. Choose column $\mathbf{Y}^i$ and consider two models α_a and α_b such that:

$$\alpha_a = (\alpha_1, \cdots, \alpha_{i-1}, 1, \alpha_{i+1}, \cdots, \alpha_p)$$
$$\alpha_b = (\alpha_1, \cdots, \alpha_{i-1}, 0, \alpha_{i+1}, \cdots, \alpha_p)$$

2. For each model, calculate B_{a0} and B_{b0} as defined by (26).
3. Sample

$$\alpha_i \mid \alpha_1, \cdots, \alpha_{i-1}, \alpha_{i+1}, \cdots, \alpha_p \sim Ber(p_i)$$

where

$$p_i = \frac{B_{a0}\mathbf{P}\left(M_{\alpha_a}\right)}{B_{a0}\mathbf{P}\left(M_{\alpha_a}\right) + B_{b0}\mathbf{P}\left(M_{\alpha_b}\right)},$$

Using this algorithm, we visit the most likely models given our set of responses. Under the model and prior specification, there are closed-form expressions for the posteriors of the model parameters β_α and σ.

A.4. Hyper Parameter for the g-prior

We use a local empirical Bayes to choose the hyper parameter for the g-prior in (25). Since we allow g to be a function of the columns of $\mathbf{Y}$ as well as the model defined by α, we calculate a separate g for each univariate Bayes factor in (25) above. An empirical Bayes estimate of g maximizes the marginal likelihood and is constrained to be non-negative. From Liang et al. (2008b), we have:

$$\hat{g}_\alpha^{EB(i)} = max\{F_\alpha^i - 1, 0\} \tag{29}$$

$$F_\alpha^i = \frac{R_\alpha^{2i}/k_\alpha}{(1 - R_\alpha^{2i})/(N - 1 - k_\alpha)}. \tag{30}$$

For univariate stochastic search, the literature recommends choosing a fixed g as the number of data points Garcia-Donato and Martinez-Beneito (2013). However, the multivariate nature of our model induced by the vector-valued response makes this approach unreliable. Since each response has distinct statistical characteristics and correlations with the covariates, it is necessary to vary g among different sampled models and responses. We find that this approach provides sufficiently stable estimation of the inclusion probabilities for the covariates.

Appendix B: Derivation of lasso form

In this section of the Appendix, we derive the penalized objective (lasso) forms of the utility functions. After integration over $p(\tilde{Y}, \tilde{X}, \Theta|\mathbf{Y}, \mathbf{X})$, the utility takes the form (from equation (10)):

$$\mathcal{L}(\boldsymbol{\gamma}) = \text{tr}[M\boldsymbol{\gamma}S\boldsymbol{\gamma}^T] - 2\text{tr}[A\boldsymbol{\gamma}^T] + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0, \tag{31}$$

where $A = \mathbb{E}[\Omega\tilde{Y}\tilde{X}^T]$, $S = \mathbb{E}[\tilde{X}\tilde{X}^T] = \overline{\Sigma_x}$, and $M = \overline{\Omega}$, and the overlines denote posterior means. Defining the Cholesky decompositions: $M = LL^T$ and $S = QQ^T$, combining the matrix traces, completing the square with respect to $\boldsymbol{\gamma}$, and converting the trace to the vectorization operator, we obtain:

$$\begin{aligned}
\mathcal{L}(\boldsymbol{\gamma}) &= \text{tr}[M(\boldsymbol{\gamma}S\boldsymbol{\gamma}^T - 2M^{-1}A\boldsymbol{\gamma}^T] + \lambda \left\|\ \mathbf{vec}(\boldsymbol{\gamma})\right\|_0 \\
&\propto \text{tr}\left[M(\boldsymbol{\gamma} - M^{-1}AS^{-1})S(\boldsymbol{\gamma} - M^{-1}AS^{-1})^T\right] + \lambda \left\|\ \mathbf{vec}(\boldsymbol{\gamma})\right\|_0 \\
&= \text{tr}\left[LL^T(\boldsymbol{\gamma} - L^{-T}L^{-1}AS^{-1})S(\boldsymbol{\gamma} - L^{-T}L^{-1}AS^{-1})^T\right] + \lambda \left\|\ \mathbf{vec}(\boldsymbol{\gamma})\right\|_0 \\
&= \text{tr}\left[L^T(\boldsymbol{\gamma} - L^{-T}L^{-1}AS^{-1})S(\boldsymbol{\gamma} - L^{-T}L^{-1}AS^{-1})^TL\right] + \lambda \left\|\ \mathbf{vec}(\boldsymbol{\gamma})\right\|_0 \\
&= \text{tr}\left[(L^T\boldsymbol{\gamma} - L^{-1}AQ^{-T}Q^{-1})QQ^T((L^T\boldsymbol{\gamma} - L^{-1}AQ^{-T}Q^{-1})^T\right] + \lambda \left\|\ \mathbf{vec}(\boldsymbol{\gamma})\right\|_0 \\
&= \text{tr}\left[(L^T\boldsymbol{\gamma}Q - L^{-1}AQ^{-T})(L^T\boldsymbol{\gamma}Q - L^{-1}AQ^{-T})^T\right] + \lambda \left\|\ \mathbf{vec}(\boldsymbol{\gamma})\right\|_0 \\
&= \ \mathbf{vec}(L^T\boldsymbol{\gamma}Q - L^{-1}AQ^{-T})^T\mathbf{vec}(L^T\boldsymbol{\gamma}Q - L^{-1}AQ^{-T}) + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0.
\end{aligned} \tag{32}$$

The proportionality in line 2 is up to an additive constant with respect to the action variable, $\boldsymbol{\gamma}$. We arrive at the final utility by distributing the vectorization and rewriting the inner product as a squared ℓ_2 norm.

$$\mathcal{L}(\boldsymbol{\gamma}) = \left\|\left[Q^T \otimes L^T\right]\mathbf{vec}(\boldsymbol{\gamma}) - \mathbf{vec}(L^{-1}AQ^{-T})\right\|_2^2 + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0. \tag{33}$$

The l_0 norm penalty yields a difficult combinatorial optimization problem even for a relatively small dimensions ($pq \approx 30$). Thus, one may use an ℓ_1 norm as the most straightforward approximation to the ℓ_0 norm, yielding the loss function:

$$\mathcal{L}(\boldsymbol{\gamma}) = \left\|\left[Q^T \otimes L^T\right]\mathbf{vec}(\boldsymbol{\gamma}) - \mathbf{vec}(L^{-1}AQ^{-T})\right\|_2^2 + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_1. \tag{34}$$

Appendix C: Derivation of the loss function under fixed predictors

We devote this section to deriving an analogous loss function for multivariate regression when the predictors are assumed fixed. Notice that this is essentially an extension of Hahn and Carvalho (2015) to the multiple response case and adds to the works of Brown, Vannucci and Fearn (1998) and Wang (2010) by providing a posterior summary strategy that relies on more than just marginal quantities like posterior inclusion probabilities.

Suppose we observe N realizations of the predictor vector defining the design matrix $\mathbf{X} \in \mathbb{R}^{Nxp}$. Future realizations $\tilde{\mathbf{Y}} \in \mathbb{R}^{Nxq}$ at this fixed set of predictors are generated from a matrix normal distribution:

$$\tilde{\mathbf{Y}} \sim \text{Matrix Normal}_{N,q}\left(\mathbf{X}\boldsymbol{\gamma}^T,\ \mathbb{I}_{NxN},\ \Psi_{qxq}\right). \tag{35}$$

In this case, the optimal posterior summary $\boldsymbol{\gamma}^*$ minimizes the expected loss $\mathcal{L}_\lambda(\boldsymbol{\gamma}) = \mathbb{E}[\mathcal{L}_\lambda(\tilde{\mathbf{Y}}, \Theta, \boldsymbol{\gamma})]$. Here, the expectation is taken over the joint space of the predictive and posterior distributions: $p(\tilde{\mathbf{Y}}, \Theta|\mathbf{Y}, \mathbf{X})$ *where* $\tilde{X}$ *is now absent since we are relegated to predicting at the observed covariate matrix* $\mathbf{X}$. We define the utility function using the negative kernel of distribution (35) where, as before, $\boldsymbol{\gamma}$ is the summary defining the sparsified linear predictor and $\Omega = \Psi^{-1}$:

$$\mathcal{L}_\lambda(\tilde{\mathbf{Y}}, \Theta, \boldsymbol{\gamma}) = \frac{1}{2}\text{tr}\left[\Omega(\tilde{\mathbf{Y}} - \mathbf{X}\boldsymbol{\gamma}^T)^T(\tilde{\mathbf{Y}} - \mathbf{X}\boldsymbol{\gamma}^T)\right] + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0, \tag{36}$$

Expanding the inner product and dropping terms that do not involve $\boldsymbol{\gamma}$, we define the loss up to proportionality:

$$\mathcal{L}_\lambda(\tilde{\mathbf{Y}}, \Theta, \boldsymbol{\gamma}) \propto \text{tr}\left[\Omega(\boldsymbol{\gamma}\mathbf{X}^T\mathbf{X}\boldsymbol{\gamma}^T - 2\tilde{\mathbf{Y}}^T\mathbf{X}\boldsymbol{\gamma}^T)\right] + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0. \tag{37}$$

Analogous to the stochastic predictors derivation, we integrate over $(\tilde{Y}, \Theta)$ to obtain our expected loss:

$$\begin{aligned}\mathcal{L}_\lambda(\boldsymbol{\gamma}) &= \mathbb{E}[\mathcal{L}_\lambda(\tilde{\mathbf{Y}}, \Theta, \boldsymbol{\gamma})] \\ &= \text{tr}[M\boldsymbol{\gamma}S_f\boldsymbol{\gamma}^T] - 2\text{tr}[A_f\boldsymbol{\gamma}^T] + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_0.\end{aligned} \tag{38}$$

where, similar to the random predictor case, $A_f = \mathbb{E}[\Omega\tilde{\mathbf{Y}}^T\mathbf{X}]$, $S_f = \mathbf{X}^T\mathbf{X}$, $M = \overline{\Omega}$, and the overlines denote posterior means. The subscript f is used to denote quantities calculated at *fixed* design points $\mathbf{X}$. Defining the Cholesky decompositions: $M = LL^T$ and $S_f = Q_f Q_f^T$ and replacing the ℓ_0 norm with the ℓ_1 norm, this expression can be formulated in the form of a standard penalized regression problem:

$$\mathcal{L}_\lambda(\boldsymbol{\gamma}) = \left\|\left[Q_f^T \otimes L^T\right]\mathbf{vec}(\boldsymbol{\gamma}) - \mathbf{vec}(L^{-1}A_f Q_f^{-T})\right\|_2^2 + \lambda \left\|\mathbf{vec}(\boldsymbol{\gamma})\right\|_1 \tag{39}$$

with covariates $Q_f^T \otimes L^T$, "data" $L^{-1}A_f Q_f^{-T}$, and regression coefficients $\boldsymbol{\gamma}$. Accordingly, (12) can be optimized using existing software such as the `lars` R package of Efron et al. (2004).

We use loss function (39) as a point of comparison to demonstrate how incorporating covariate uncertainty may impact the summarization procedure in our applications.

References

Bayarri, M. J., Berger, J. O., Forte, A. and Garcia-Donato, G. (2012). Criteria for Bayesian model choice with application to variable selection. *The Annals of Statistics* **40** 1550–1577.

Brown, P. J., Vannucci, M. and Fearn, T. (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **60** 627–641.

Clyde, M. and George, E. I. (2004). Model uncertainty. *Statistical Science* **19** 81–94.

Cochrane, J. H. (2011). Presidential address: Discount rates. *The Journal of Finance* **66** 1047–1108.

Dobra, A., Hans, C., Jones, B., Nevins, J. R., Yao, G. and West, M. (2004). Sparse graphical models for exploring gene expression data. *Journal of Multivariate Analysis* **90** 196–212.

Efron, B., Hastie, T., Johnstone, I., Tibshirani, R. et al. (2004). Least angle regression. *The Annals of statistics* **32** 407–499.

Ericsson, J. and Karlsson, S. (2004). Choosing Factors in a Multifactor Asset Pricing Model: A Bayesian Approach Technical Report, Stockholm School of Economics.

Fama, E. F. and French, K. R. (1992). The cross-section of expected stock returns. *the Journal of Finance* **47** 427–465.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics* **116** 1–22.

Garcia-Donato, G. and Martinez-Beneito, M. (2013). On sampling strategies in Bayesian variable selection problems with large model spaces. *Journal of the American Statistical Association* **108** 340–352.

George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association* **88** 881–889.

Hahn, P. R. and Carvalho, C. M. (2015). Decoupling shrinkage and selection in Bayesian linear models: a posterior summary perspective. *Journal of the American Statistical Association* **110** 435–448.

Hans, C., Dobra, A. and West, M. (2007). Shotgun stochastic search for "large p" regression. *Journal of the American Statistical Association* **102** 507–516.

Harvey, C. R. and Liu, Y. (2015). Lucky factors. *Available at SSRN 2528780*.

Jeffreys, H. (1961). Theory of Probability (3rd edt.) Oxford University Press.

Jones, B., Carvalho, C., Dobra, A., Hans, C., Carter, C. and West, M. (2005). Experiments in stochastic computation for high-dimensional graphical models. *Statistical Science* 388–400.

Liang, F., Paulo, R., Molina, G., Clyde, M. and Berger, J. (2008a). Mixtures of g Priors for Bayesian Variable Selection. *Journal of the American Statistical Association* **103** 410–423.

Liang, F., Paulo, R., Molina, G., Clyde, M. A. and Berger, J. O. (2008b). Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association* **103**.

Murray, J. (2015). bfa.

Murray, J. S., Dunson, D. B., Carin, L. and Lucas, J. E. (2013). Bayesian Gaussian copula factor models for mixed data. *Journal of the American Statistical Association* **108** 656–665.

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of economic theory* **13** 341–360.

Scott, J. and Berger, J. O. (2006). An exploration of aspects of Bayesian multiple testing. *Journal of Statistical Planning and Inference* **136** 2144–2162.

Wang, H. (2010). Sparse seemingly unrelated regression modelling: Applications

in finance and econometrics. *Computational Statistics & Data Analysis* **54** 2866–2877.

Wang, H. and West, M. (2009). Bayesian analysis of matrix normal graphical models. *Biometrika* **96** 821–834.

Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American statistical Association* **57** 348–368.

Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. *Bayesian inference and decision techniques: Essays in Honor of Bruno De Finetti* **6** 233–243.

Zellner, A. and Siow, A. (1980). Posterior odds ratios for selected regression hypotheses. *Trabajos de estadística y de investigación operativa* **31** 585–603.

Zellner, A. and Siow, A. (1984). *Basic issues in econometrics*. University of Chicago Press Chicago.